-Merging of Linux/uClinux 2.6
& the Benchmark-

CE Linux Forum

Korea Tech Conference

2005년 5월 14일, 서울

Merging of Linux/uClinux 2.6 & the Benchmark

Hyok S. Choi (최 혁 승)
Linux Kernel armnommu maintainer

Digital Media R&D Center
Samsung Electronics Co.,Ltd.

Contents

- Introduction of uClinux
- Introduction of
  Linux 2.6 for MMU-less ARM Project
- Recent Changes of ARM Linux Kernel
- The Benchmark
- What's the next?

Introduction of uClinux(1/2)

- **What is uClinux?**
  - A Linux derivative which is independent from the H/W supported Paging Management of MMU.
  - The first uClinux - 1998, Linux 2.0
  - Currently, under merging state into the mainline kernel 2.6. (m68knommu, v850, h8300 is done)
  - Supported Architectures :
    - Motorola M68K/ColdFire, ARM 7/9/10/11, Intel i960, Sun SPARC, ADI BlackFin, Axis Etrax, PRISMA, Atari 68k, Xilinx Microblaze, NEC v850, Hitachi H8
  - Market and Devices :
    - Gateways, VoIP phones, Blutooth devices, web-cams, Auto Vehicle Locators, Security Appliances, Handhelds

Introduction of uClinux(2/2)

**“The one of
the most used Linux distribution
in real embedded systems
on commercial product.”**

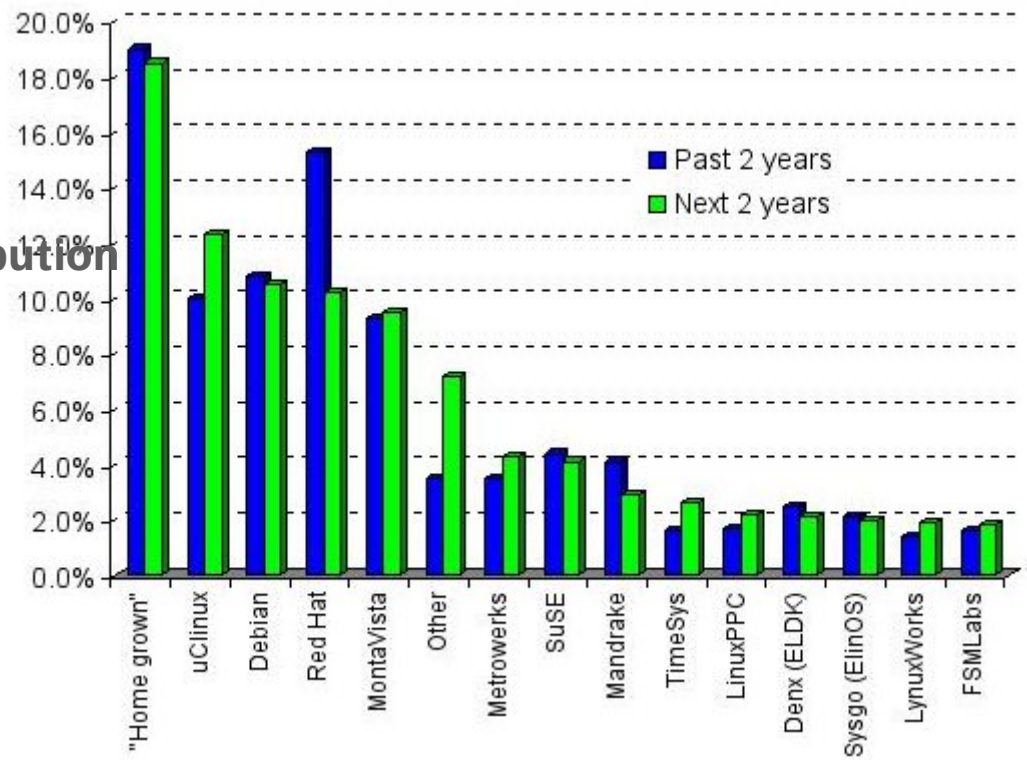

- Snapshot of the Embedded Linux market -- March, 2004 , linuxdevices.com

Introduction of
Linux 2.6 for MMU-less ARM Project (1/3)

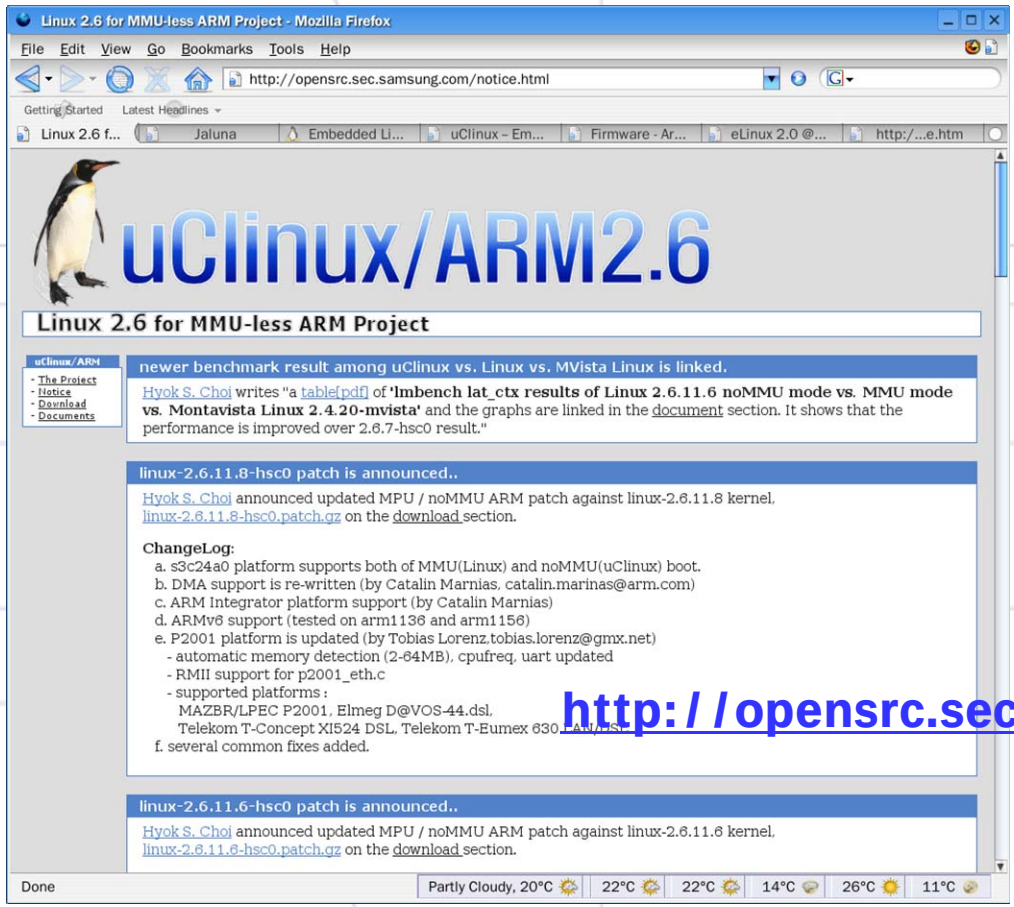

http://opensrc.sec.samsung.com/

Introduction of Linux 2.6 for MMU-less ARM Project (2/3)

- **Latest version : 2.6.12-rc3-mm3-hsc0**
- **URL : http://opensrc.sec.samsung.com/**
- **Supported Platforms:**
  - ARM7
    - Atmel AT89x(7tdmi), Samsung S3C3410(7tdmi), S3C4510b(7tdmi), S3C44B0X(7tdmi)
  - ARM9
    - Samsung S3C24A0(926ej), S5C7375(920T), S5H5002(940T), P2001(9TDMI)
  - ARM11
    - ARM Integrator-CP Series(various including v6 architectures)
  - Known to Support
    - TI DM270, Philips LPC22xx, Apple iPod, S3C2500

    * Sony Clie-SL10, Nintendo-DS

Introduction of Linux 2.6 for MMU-less ARM Project (3/3)

- **Why uClinux/ARM 2.6?**
  - Lightweight
    - Complete Linux 2.6 preemptible kernel zImage : 300KB
    - 30~50% lighter application binary (flat-binary/uClibc/c++)
  - Light Latency
    - Much faster context switching, FIFO troughput
  - XIP (eXecute In Place)
  - Cheaper
    - MMU core size in typical ARM SoC wafer is about 30%.
  - Full Linux API
    - Support the full Linux API, with few exceptions like fork().
  - Easier to adapt
    - Firmware code runs as a plain uClinux application with a bit of wrapper.
  - Full Linux 2.6 kernel features
    - Supports full filesystems, device drivers.

Recent Changes of ARM Linux Kernel

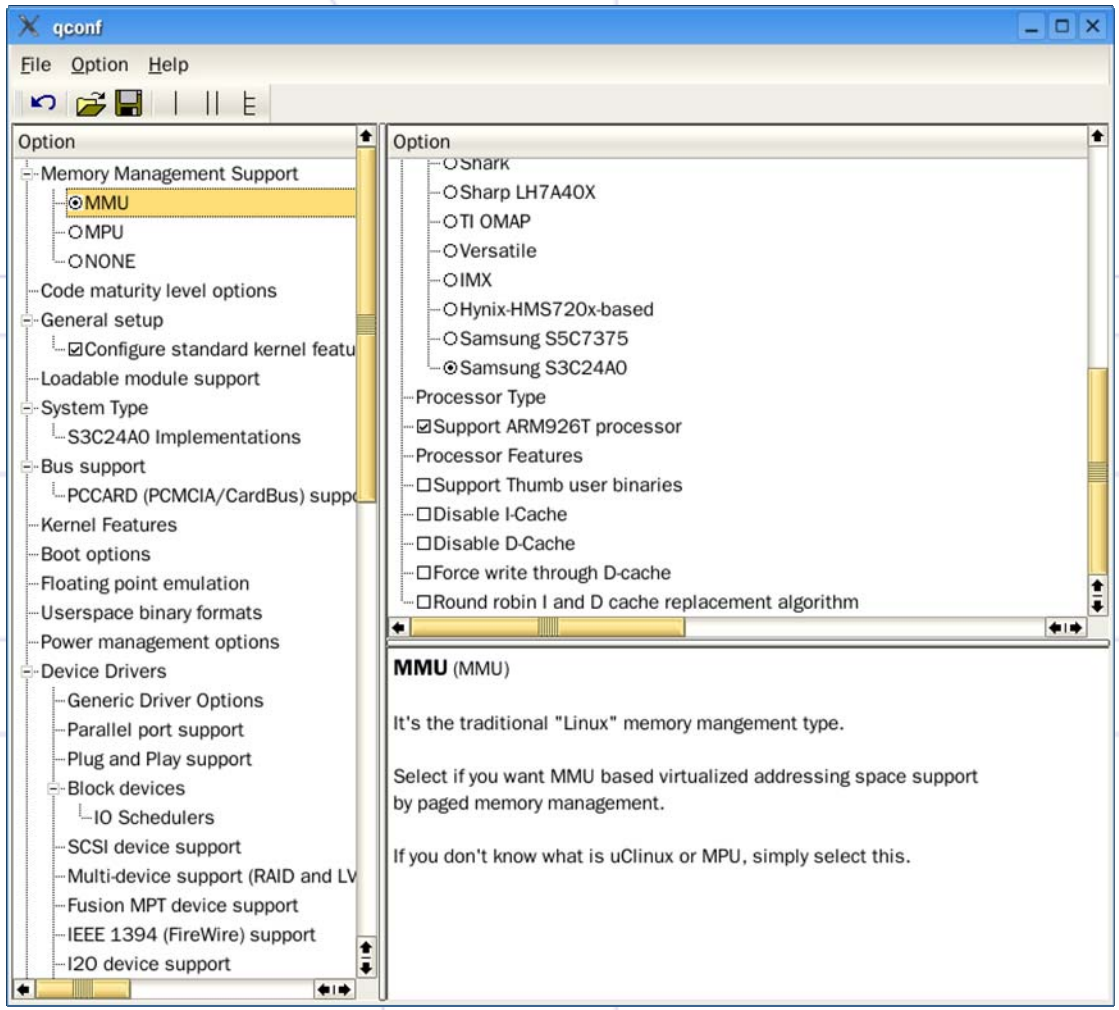

Recent Changes of ARM Linux Kernel

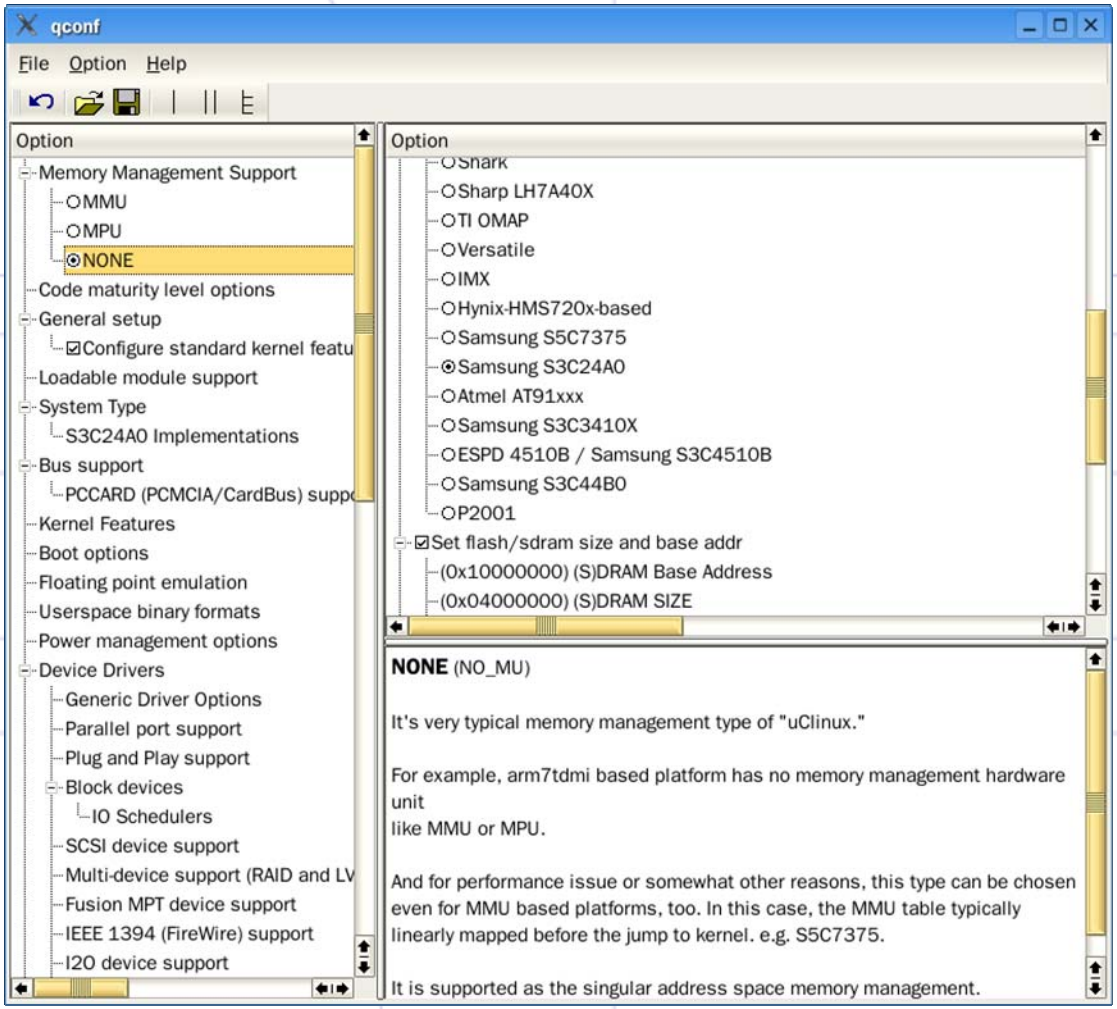

Recent Changes of ARM Linux Kernel

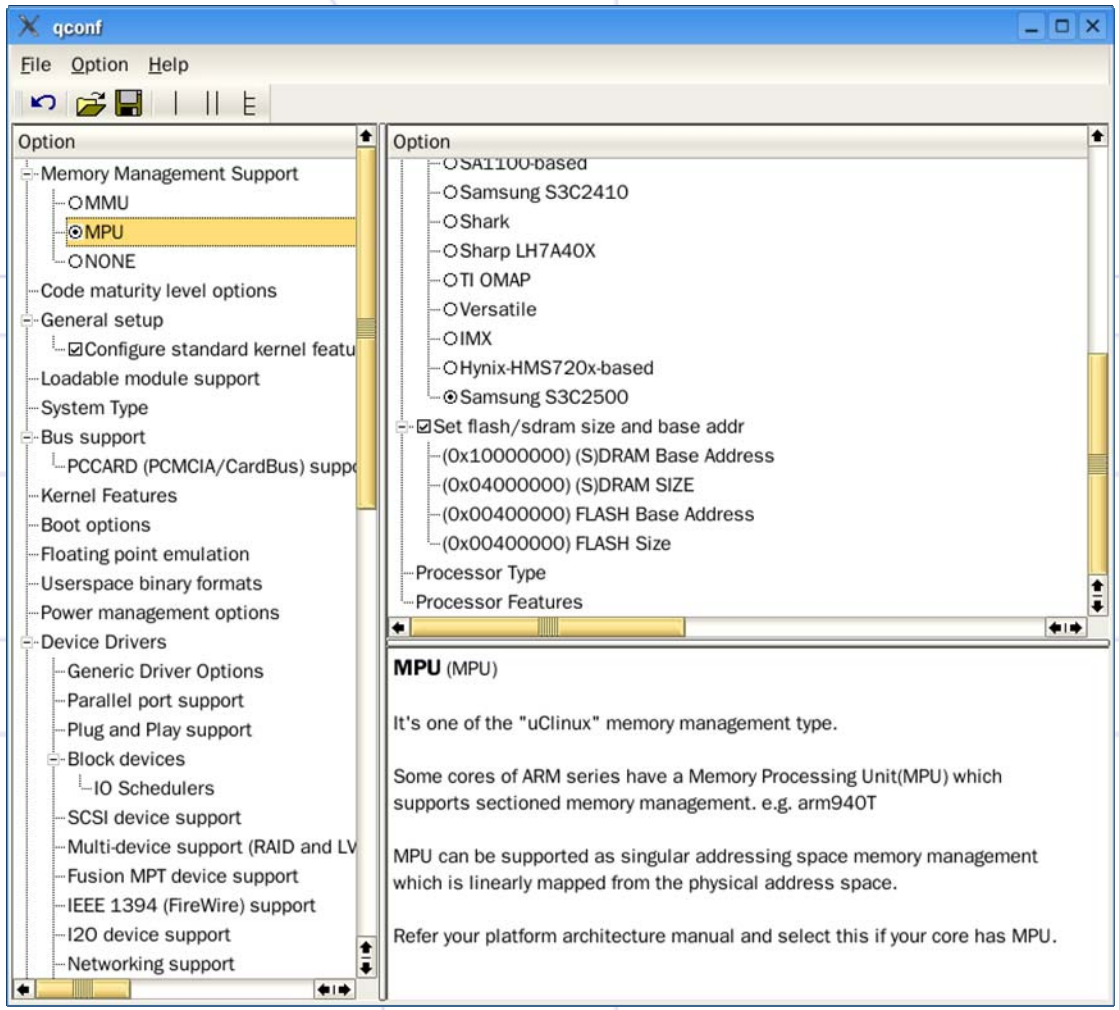

Recent Changes of ARM Linux Kernel

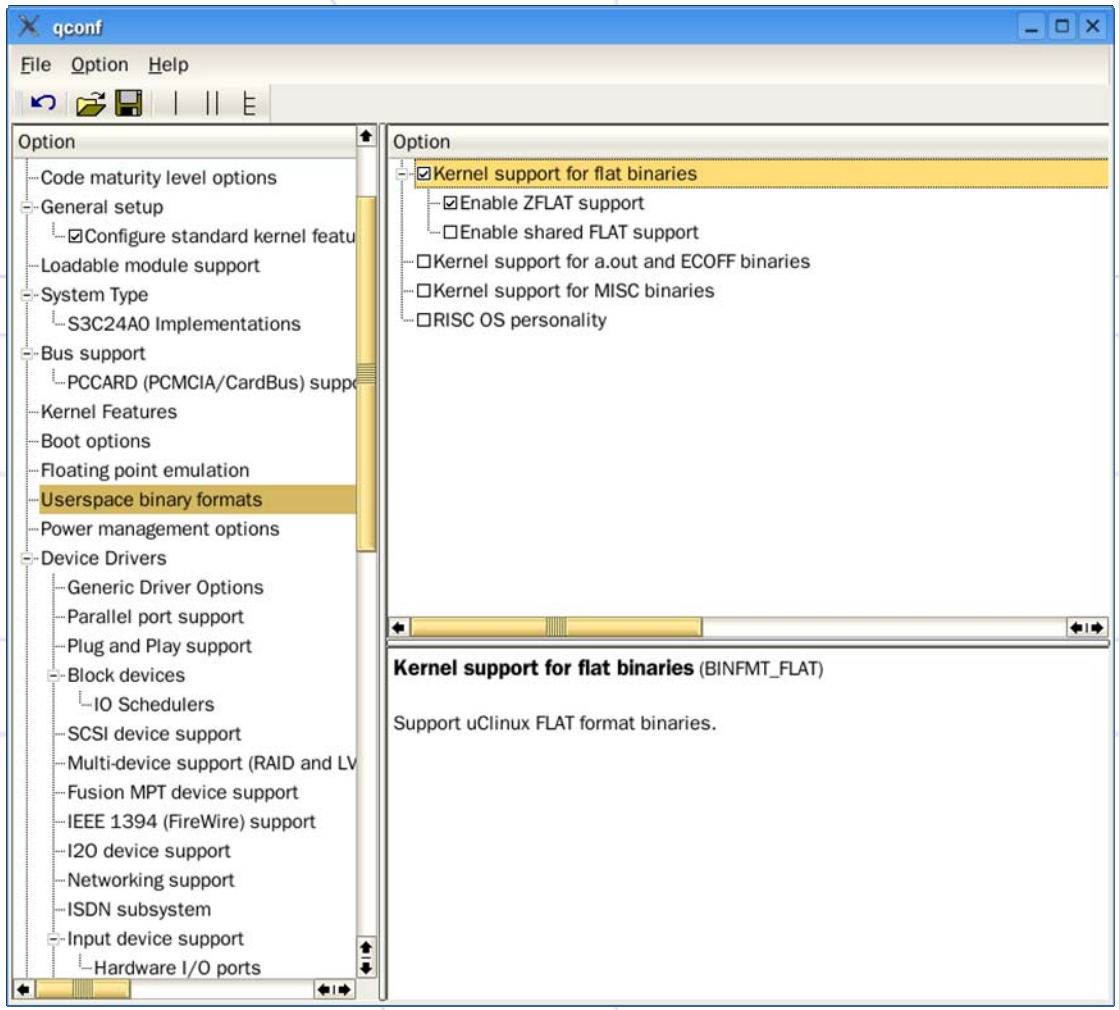

Recent Changes of ARM Linux Kernel

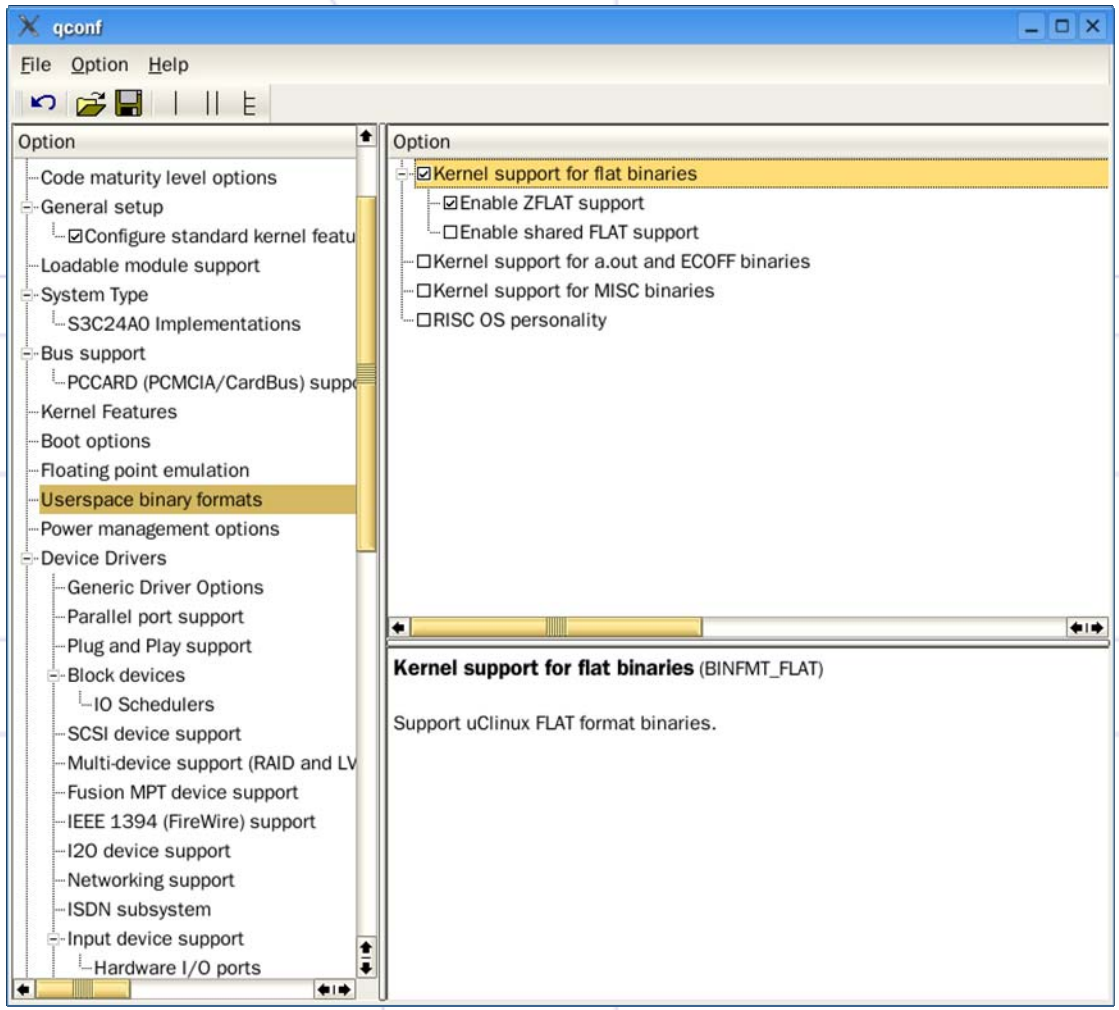

Recent Changes of ARM Linux Kernel

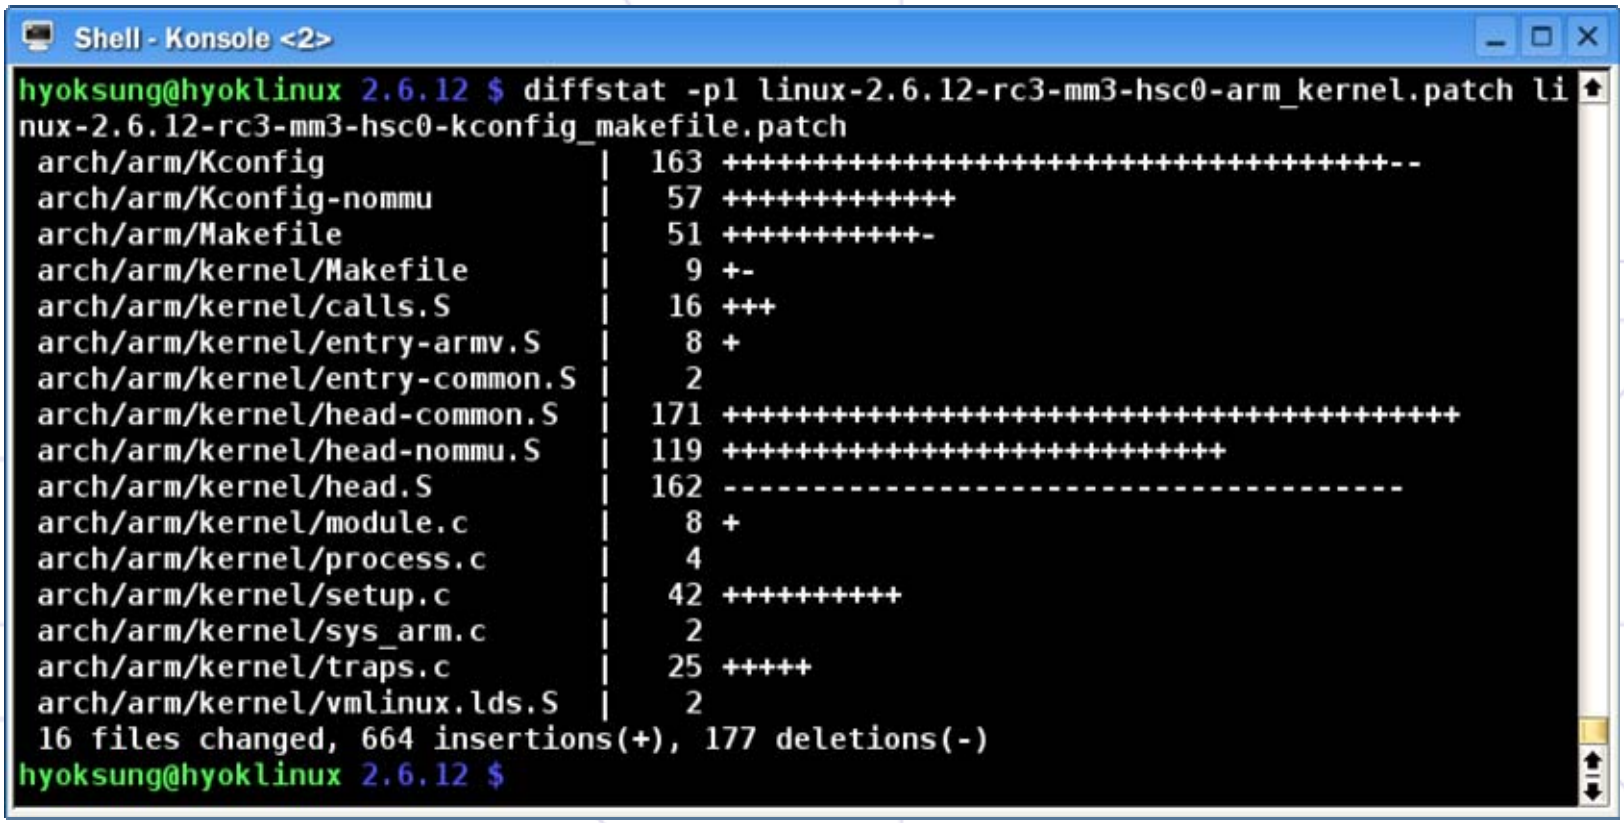

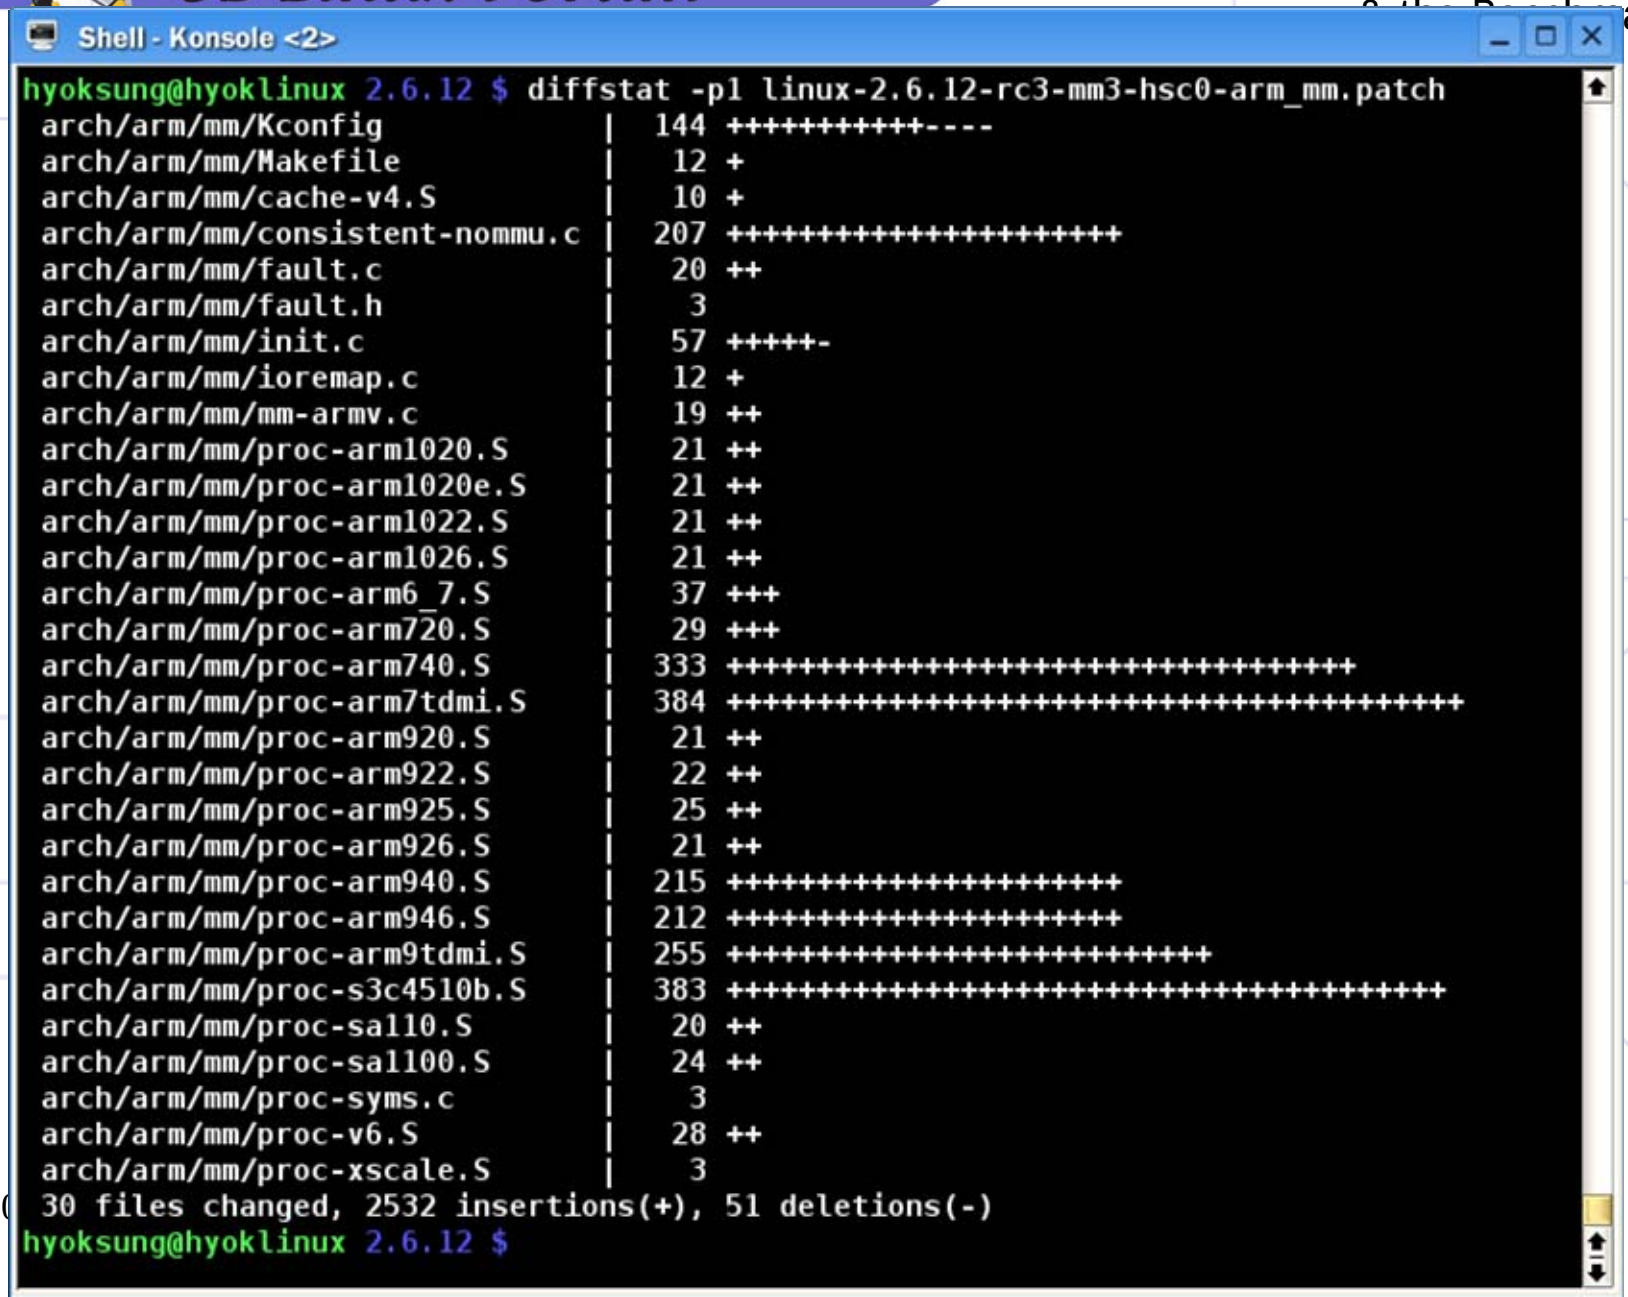
Shell - Konsole <2>
hyoksung@hyoklinux 2.6.12 $ diffstat -p1 linux-2.6.12-rc3-mm3-hsc0-arm_mm.patch
arch/arm/mm/Kconfig | 144 ++++++++++----
arch/arm/mm/Makefile | 12 +
arch/arm/mm/cache-v4.S | 10 +
arch/arm/mm/consistent-nommu.c | 207 ++++++++++++++++++++
arch/arm/mm/fault.c | 20 ++
arch/arm/mm/fault.h | 3
arch/arm/mm/init.c | 57 +++++-
arch/arm/mm/ioremap.c | 12 +
arch/arm/mm/mm-armv.c | 19 ++
arch/arm/mm/proc-arm1020.S | 21 ++
arch/arm/mm/proc-arm1020e.S | 21 ++
arch/arm/mm/proc-arm1022.S | 21 ++
arch/arm/mm/proc-arm1026.S | 21 ++
arch/arm/mm/proc-arm6_7.S | 37 +++
arch/arm/mm/proc-arm720.S | 29 +++
arch/arm/mm/proc-arm740.S | 333 ++++++++++++++++++++++++++++++++++
arch/arm/mm/proc-arm7tdmi.S | 384 ++++++++++++++++++++++++++++++++++++++
arch/arm/mm/proc-arm920.S | 21 ++
arch/arm/mm/proc-arm922.S | 22 ++
arch/arm/mm/proc-arm925.S | 25 ++
arch/arm/mm/proc-arm926.S | 21 ++
arch/arm/mm/proc-arm940.S | 215 ++++++++++++++++++++
arch/arm/mm/proc-arm946.S | 212 ++++++++++++++++++++
arch/arm/mm/proc-arm9tdmi.S | 255 ++++++++++++++++++++++
arch/arm/mm/proc-s3c4510b.S | 383 ++++++++++++++++++++++++++++++++++++++
arch/arm/mm/proc-sa110.S | 20 ++
arch/arm/mm/proc-sa1100.S | 24 ++
arch/arm/mm/proc-syms.c | 3
arch/arm/mm/proc-v6.S | 28 ++
arch/arm/mm/proc-xscale.S | 3
30 files changed, 2532 insertions(+), 51 deletions(-)
hyoksung@hyoklinux 2.6.12 $

-Merging of Linux/uClinux 2.6
ark-

Shell - Konsole <2>

```
hyoksung@hyoklinux 2.6.12 $ diffstat -p1 linux-2.6.12-rc3-mm3-hsc0-asm.patch
 include/asm-arm/bugs.h                |    5 +
 include/asm-arm/byteorder.h           |    8 ++
 include/asm-arm/cacheflush-nommu.h    |   49 ++++++++++++++++++
 include/asm-arm/cacheflush.h          |   30 ++++++++++
 include/asm-arm/cpu-multi32.h         |    5 +
 include/asm-arm/cpu-single.h          |    5 +
 include/asm-arm/domain.h              |    6 ++
 include/asm-arm/flat.h                |   20 ++++++
 include/asm-arm/glue.h                |   24 ++++++++
 include/asm-arm/hardware.h            |   23 +++++++
 include/asm-arm/hardware/dcc.h        |   49 ++++++++++++++++++
 include/asm-arm/mach/arch.h           |   15 +++++
 include/asm-arm/memory.h              |   33 ++++++++++
 include/asm-arm/mmu.h                 |    8 ++
 include/asm-arm/mmu_context.h         |    6 ++
 include/asm-arm/nommu.h               |   19 ++++++
 include/asm-arm/nommu_context.h       |   46 ++++++++++++++++
 include/asm-arm/page-nommu.h          |   54 ++++++++++++++++++++
 include/asm-arm/page.h                |    6 ++
 include/asm-arm/pgalloc.h             |    7 ++
 include/asm-arm/pgtable-nommu.h       |  108 ++++++++++++++++++++++++++++++++++++
 include/asm-arm/pgtable.h             |    7 ++
 include/asm-arm/proc-fns.h            |   35 ++++++++++++
 include/asm-arm/processor.h           |    8 ++
 include/asm-arm/procinfo.h            |   11 +++
 include/asm-arm/system.h              |    6 ++
 include/asm-arm/tlb.h                 |   10 +++
 include/asm-arm/tlbflush.h            |    9 +++
 include/asm-arm/uaccess-nommu.h       |   37 ++++++++++++
 include/asm-arm/uaccess.h             |   10 +++
 30 files changed, 659 insertions(+)
hyoksung@hyoklinux 2.6.12 $
```

Recent Changes of ARM Linux Kernel

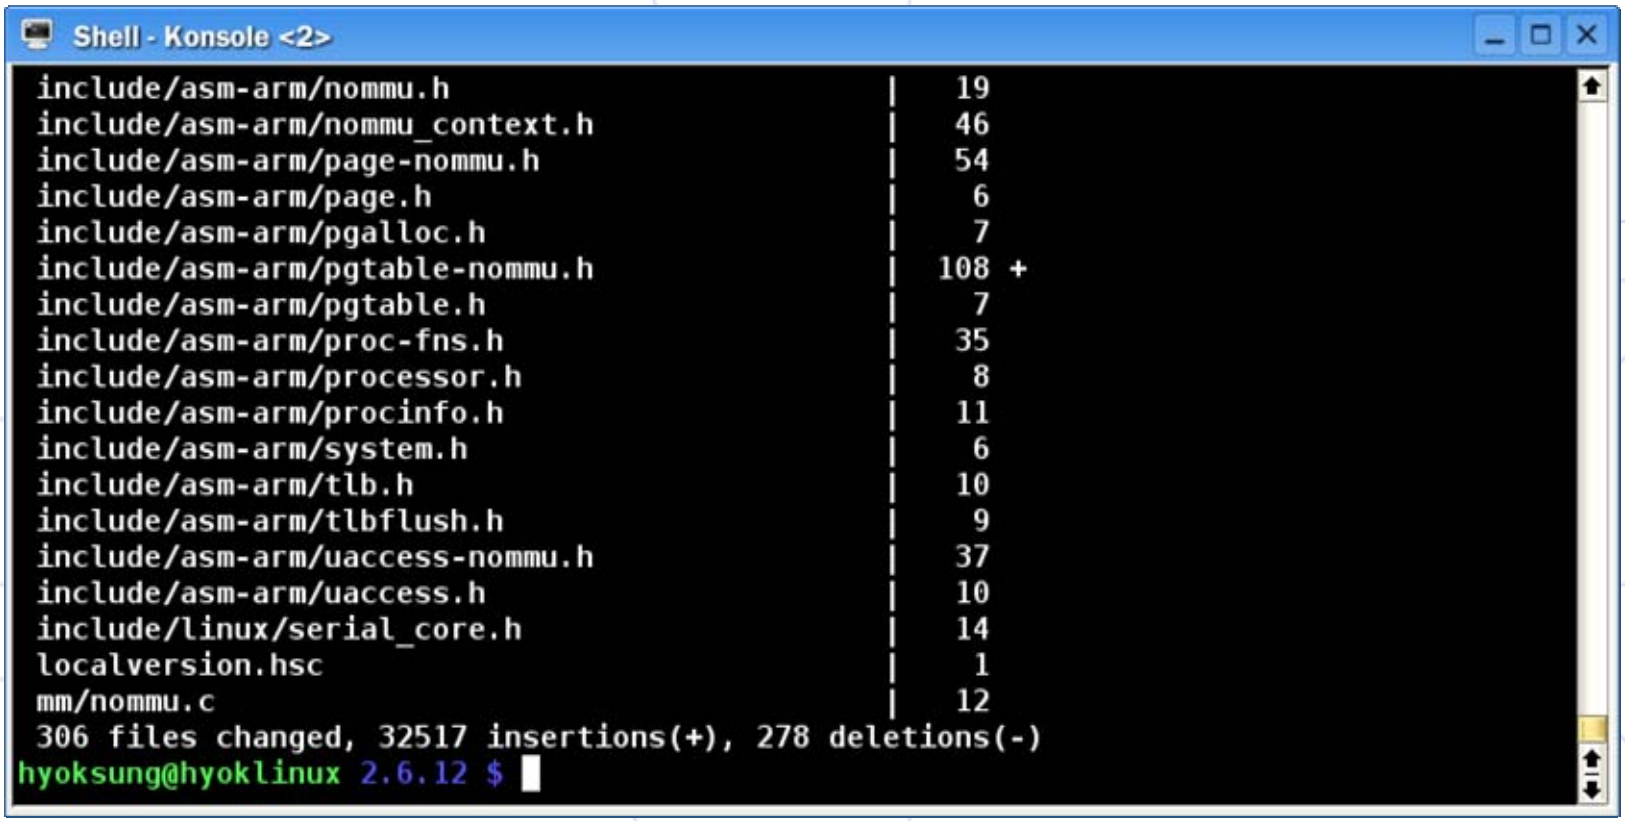

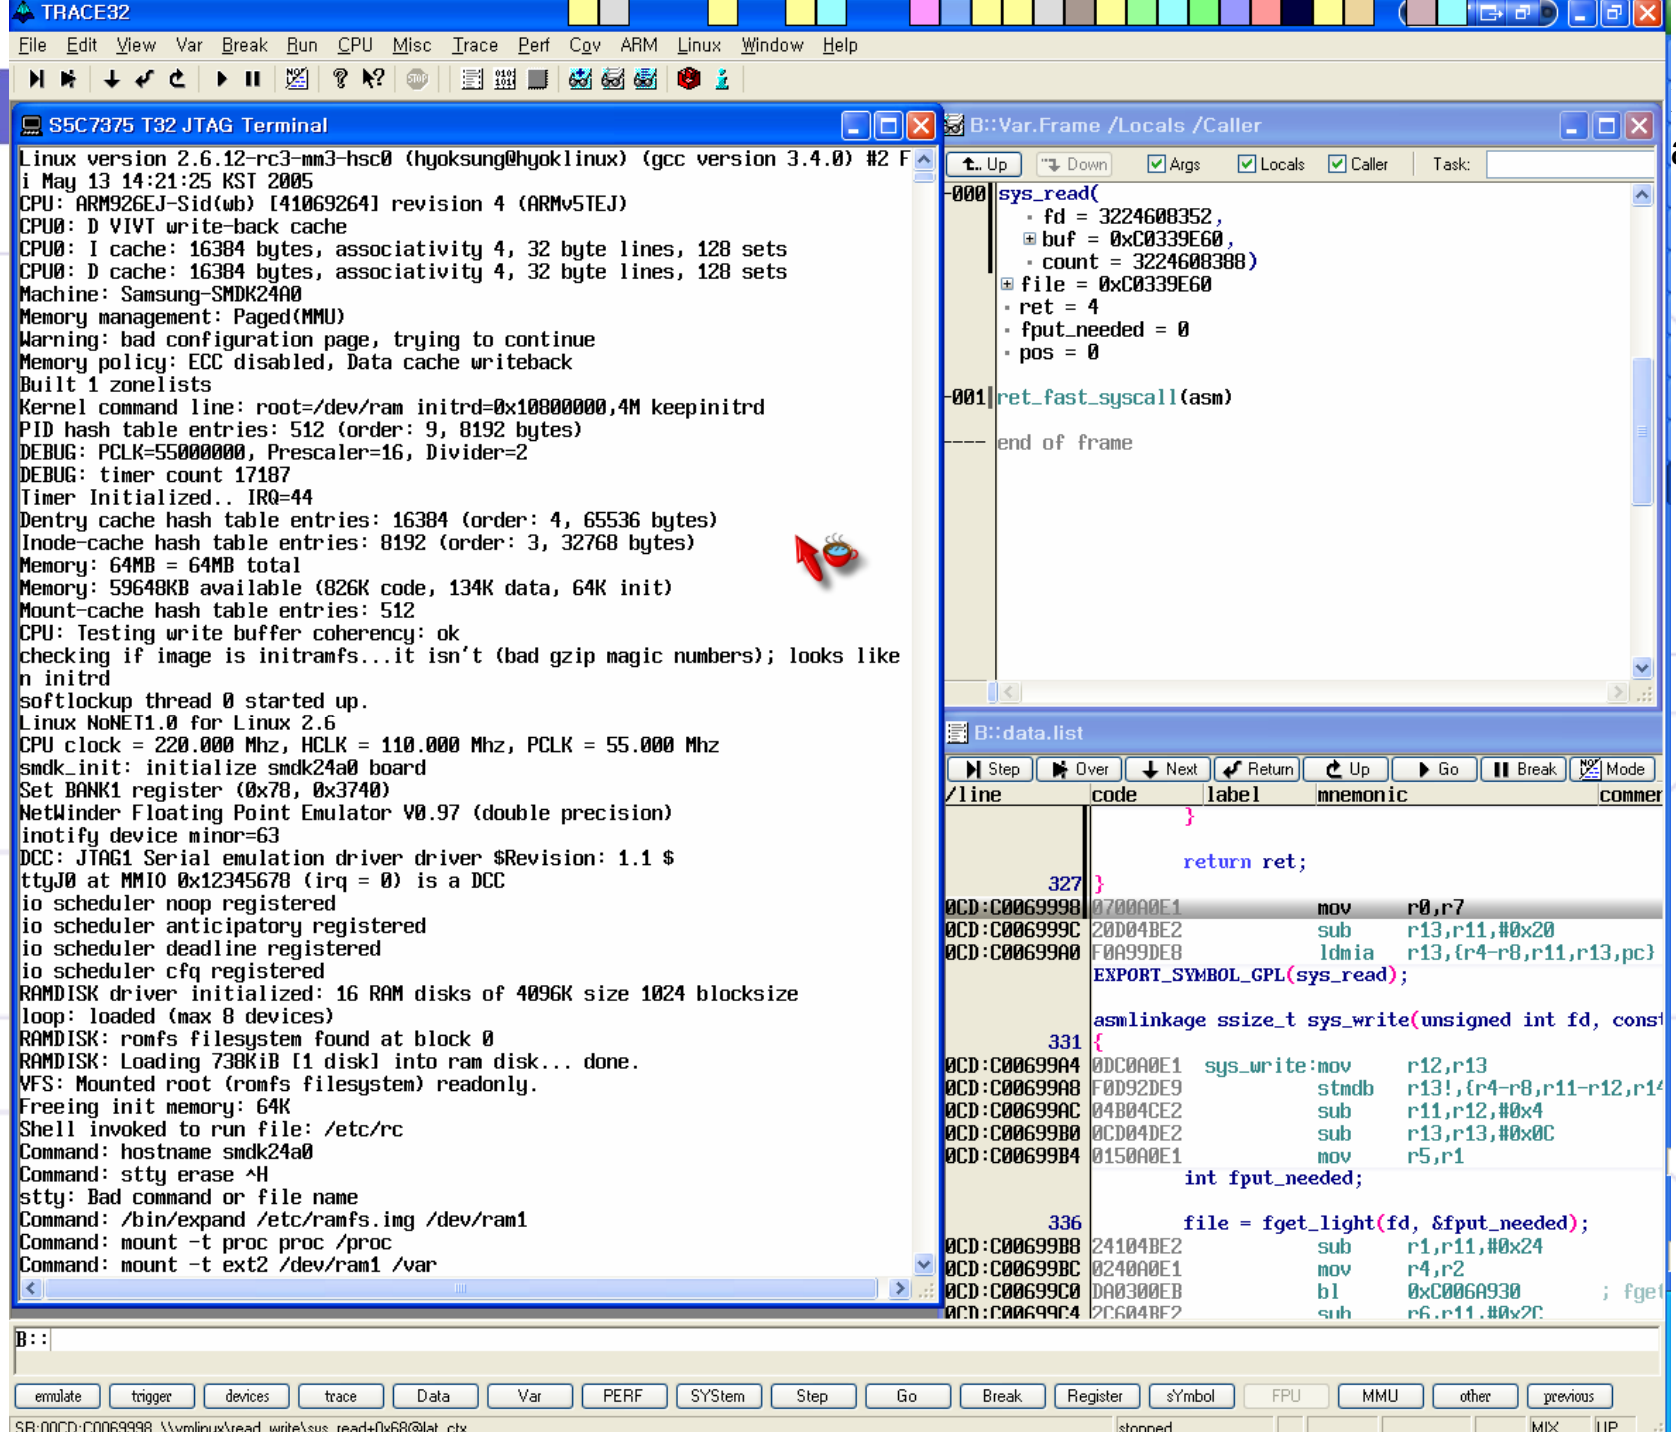
TRACE32
File Edit View Var Break Run CPU Misc Trace Perf Cov ARM Linux Window Help
S5C7375 T32 JTAG Terminal
Linux version 2.6.12-rc3-mm3-hsc0 (hyoksung@hyoklinux) (gcc version 3.4.0) #2 F
i May 13 14:21:25 KST 2005
CPU: ARM926EJ-Sid(wb) [41069264] revision 4 (ARMv5TEJ)
CPU0: D VIVT write-back cache
CPU0: I cache: 16384 bytes, associativity 4, 32 byte lines, 128 sets
CPU0: D cache: 16384 bytes, associativity 4, 32 byte lines, 128 sets
Machine: Samsung-SMDK24A0
Memory management: Paged(MMU)
Warning: bad configuration page, trying to continue
Memory policy: ECC disabled, Data cache writeback
Built 1 zonelists
Kernel command line: root=/dev/ram initrd=0x10800000,4M keepinitrd
PID hash table entries: 512 (order: 9, 8192 bytes)
DEBUG: PCLK=55000000, Prescaler=16, Divider=2
DEBUG: timer count 17187
Timer Initialized.. IRQ=44
Dentry cache hash table entries: 16384 (order: 4, 65536 bytes)
Inode-cache hash table entries: 8192 (order: 3, 32768 bytes)
Memory: 64MB = 64MB total
Memory: 59648KB available (826K code, 134K data, 64K init)
Mount-cache hash table entries: 512
CPU: Testing write buffer coherency: ok
checking if image is initramfs...it isn't (bad gzip magic numbers); looks like
n initrd
softlockup thread 0 started up.
Linux NoNET1.0 for Linux 2.6
CPU clock = 220.000 Mhz, HCLK = 110.000 Mhz, PCLK = 55.000 Mhz
smdk_init: initialize smdk24a0 board
Set BANK1 register (0x78, 0x3740)
NetWinder Floating Point Emulator V0.97 (double precision)
inotify device minor=63
DCC: JTAG1 Serial emulation driver driver $Revision: 1.1 $
ttyJ0 at MMIO 0x12345678 (irq = 0) is a DCC
io scheduler noop registered
io scheduler anticipatory registered
io scheduler deadline registered
io scheduler cfq registered
RAMDISK driver initialized: 16 RAM disks of 4096K size 1024 blocksize
loop: loaded (max 8 devices)
RAMDISK: romfs filesystem found at block 0
RAMDISK: Loading 738KiB [1 disk] into ram disk... done.
VFS: Mounted root (romfs filesystem) readonly.
Freeing init memory: 64K
Shell invoked to run file: /etc/rc
Command: hostname smdk24a0
Command: stty erase ^H
stty: Bad command or file name
Command: /bin/expand /etc/ramfs.img /dev/ram1
Command: mount -t proc proc /proc
Command: mount -t ext2 /dev/ram1 /var
B::Var.Frame /Locals /Caller
Up Down Args Locals Caller Task:
-000 sys_read(
fd = 3224608352,
buf = 0xC0339E60,
count = 3224608388)
file = 0xC0339E60
ret = 4
fput_needed = 0
pos = 0
-001 ret_fast_syscall(asm)
---- end of frame
B::data.list
Step Over Next Return Up Go Break Mode
/line code label mnemonic commer
}
return ret;
327 }
0CD:C0069998 0700A0E1 mov r0,r7
0CD:C006999C 20D04BE2 sub r13,r11,#0x20
0CD:C00699A0 F0A99DE8 ldmia r13,{r4-r8,r11,r13,pc}
EXPORT_SYMBOL_GPL(sys_read);
asmlinkage ssize_t sys_write(unsigned int fd, cons
331 {
0CD:C00699A4 0DC0A0E1 sys_write:mov r12,r13
0CD:C00699A8 F0D92DE9 stmdb r13!,{r4-r8,r11-r12,r1
0CD:C00699AC 04B04CE2 sub r11,r12,#0x4
0CD:C00699B0 0CD04DE2 sub r13,r13,#0x0C
0CD:C00699B4 0150A0E1 mov r5,r1
int fput_needed;
336 file = fget_light(fd, &fput_needed);
0CD:C00699B8 24104BE2 sub r1,r11,#0x24
0CD:C00699BC 0240A0E1 mov r4,r2
0CD:C00699C0 DA0300EB bl 0xC006A930 ; fget
B::
emulate trigger devices trace Data Var PERF SYStem Step Go Break Register sYmbol FPU MMU other previous
stopped

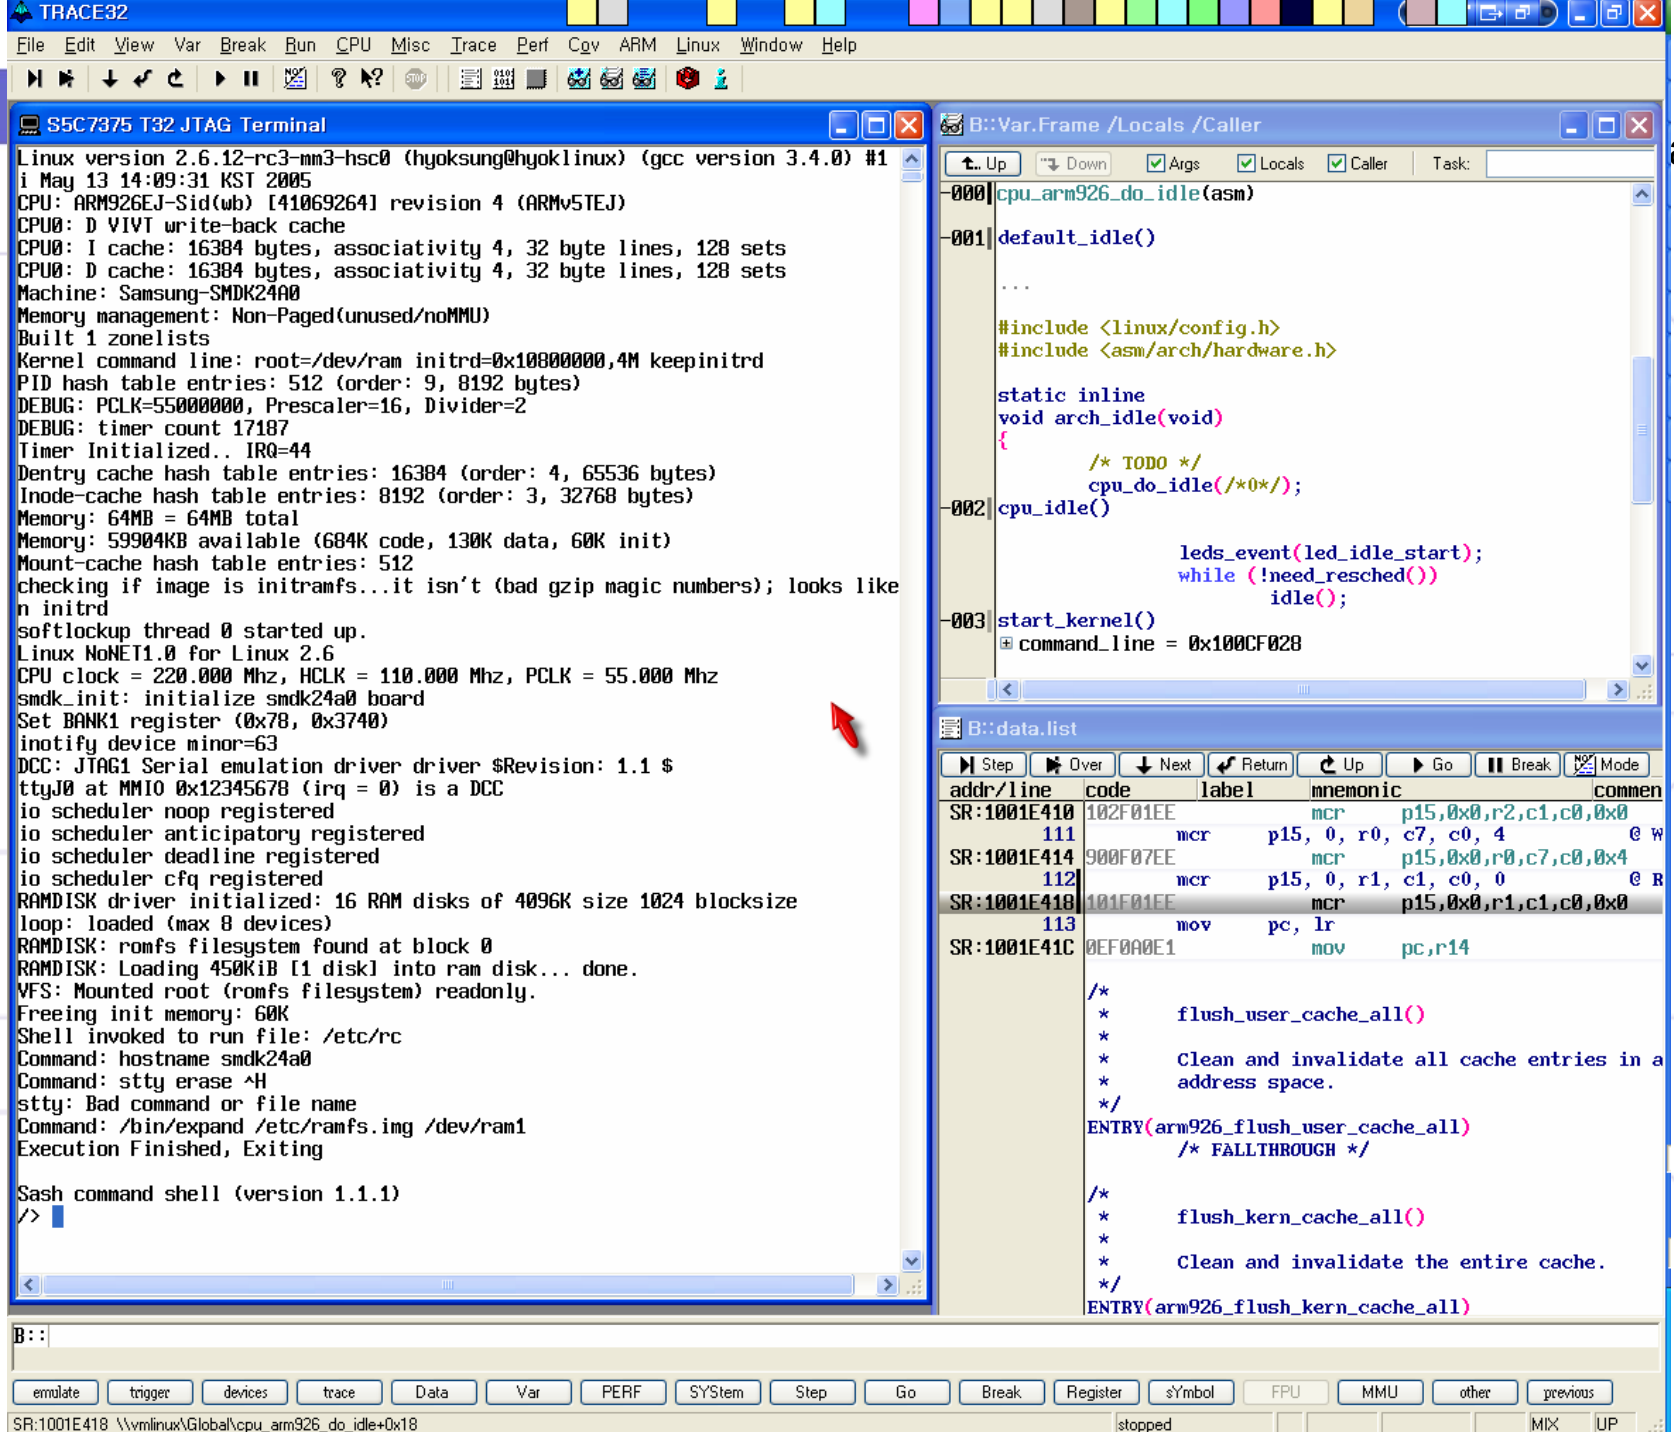
TRACE32
File Edit View Var Break Run CPU Misc Trace Perf Cov ARM Linux Window Help
S5C7375 T32 JTAG Terminal
Linux version 2.6.12-rc3-mm3-hsc0 (hyoksung@hyoklinux) (gcc version 3.4.0) #1
i May 13 14:09:31 KST 2005
CPU: ARM926EJ-Sid(wb) [41069264] revision 4 (ARMv5TEJ)
CPU0: D VIVT write-back cache
CPU0: I cache: 16384 bytes, associativity 4, 32 byte lines, 128 sets
CPU0: D cache: 16384 bytes, associativity 4, 32 byte lines, 128 sets
Machine: Samsung-SMDK24A0
Memory management: Non-Paged(unused/noMMU)
Built 1 zonelists
Kernel command line: root=/dev/ram initrd=0x10800000,4M keepinitrd
PID hash table entries: 512 (order: 9, 8192 bytes)
DEBUG: PCLK=55000000, Prescaler=16, Divider=2
DEBUG: timer count 17187
Timer Initialized.. IRQ=44
Dentry cache hash table entries: 16384 (order: 4, 65536 bytes)
Inode-cache hash table entries: 8192 (order: 3, 32768 bytes)
Memory: 64MB = 64MB total
Memory: 59904KB available (684K code, 130K data, 60K init)
Mount-cache hash table entries: 512
checking if image is initramfs...it isn't (bad gzip magic numbers); looks like
n initrd
softlockup thread 0 started up.
Linux NoNET1.0 for Linux 2.6
CPU clock = 220.000 Mhz, HCLK = 110.000 Mhz, PCLK = 55.000 Mhz
smdk_init: initialize smdk24a0 board
Set BANK1 register (0x78, 0x3740)
inotify device minor=63
DCC: JTAG1 Serial emulation driver driver $Revision: 1.1 $
ttyJ0 at MMIO 0x12345678 (irq = 0) is a DCC
io scheduler noop registered
io scheduler anticipatory registered
io scheduler deadline registered
io scheduler cfq registered
RAMDISK driver initialized: 16 RAM disks of 4096K size 1024 blocksize
loop: loaded (max 8 devices)
RAMDISK: romfs filesystem found at block 0
RAMDISK: Loading 450KiB [1 disk] into ram disk... done.
VFS: Mounted root (romfs filesystem) readonly.
Freeing init memory: 60K
Shell invoked to run file: /etc/rc
Command: hostname smdk24a0
Command: stty erase ^H
stty: Bad command or file name
Command: /bin/expand /etc/ramfs.img /dev/ram1
Execution Finished, Exiting
Sash command shell (version 1.1.1)
/>
B::Var.Frame /Locals /Caller
Up Down Args Locals Caller Task:
-000 cpu_arm926_do_idle(asm)
-001 default_idle()
...
#include <linux/config.h>
#include <asm/arch/hardware.h>
static inline
void arch_idle(void)
{
/* TODO */
cpu_do_idle(/*0*/);
-002 cpu_idle()
leds_event(led_idle_start);
while (!need_resched())
idle();
-003 start_kernel()
command_line = 0x100CF028
B::data.list
Step Over Next Return Up Go Break Mode
addr/line code label mnemonic commen
SR:1001E410 102F01EE mcr p15,0x0,r2,c1,c0,0x0
111 mcr p15, 0, r0, c7, c0, 4 @ W
SR:1001E414 900F07EE mcr p15,0x0,r0,c7,c0,0x4
112 mcr p15, 0, r1, c1, c0, 0 @ R
SR:1001E418 101F01EE mcr p15,0x0,r1,c1,c0,0x0
113 mov pc, lr
SR:1001E41C 0EF0A0E1 mov pc,r14
/*
* flush_user_cache_all()
*
* Clean and invalidate all cache entries in a
* address space.
*/
ENTRY(arm926_flush_user_cache_all)
/* FALLTHROUGH */
/*
* flush_kern_cache_all()
*
* Clean and invalidate the entire cache.
*/
ENTRY(arm926_flush_kern_cache_all)
B::
emulate trigger devices trace Data Var PERF SYStem Step Go Break Register sYmbol FPU MMU other previous
SR:1001E418 \\vmlinux\Global\cpu_arm926_do_idle+0x18
stopped
MIX UP

2.6
ark-

The Benchmark (1/6)

- **ARM9 Cache and TLB architecture**

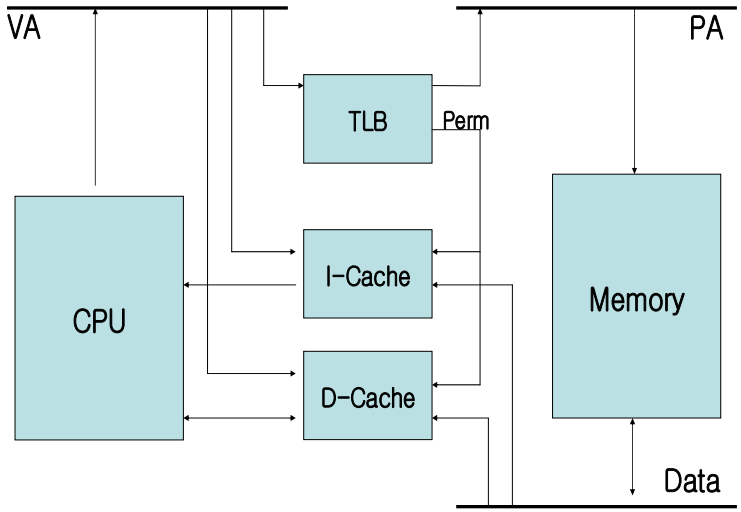

The Benchmark (2/6)

- **Imbench – benchmark program for performance testing over UNIX (McVoy. L., Staelin. C., USENIX Proceedings 1996)**
- **The FIFO structure of the modified lat_ctx**

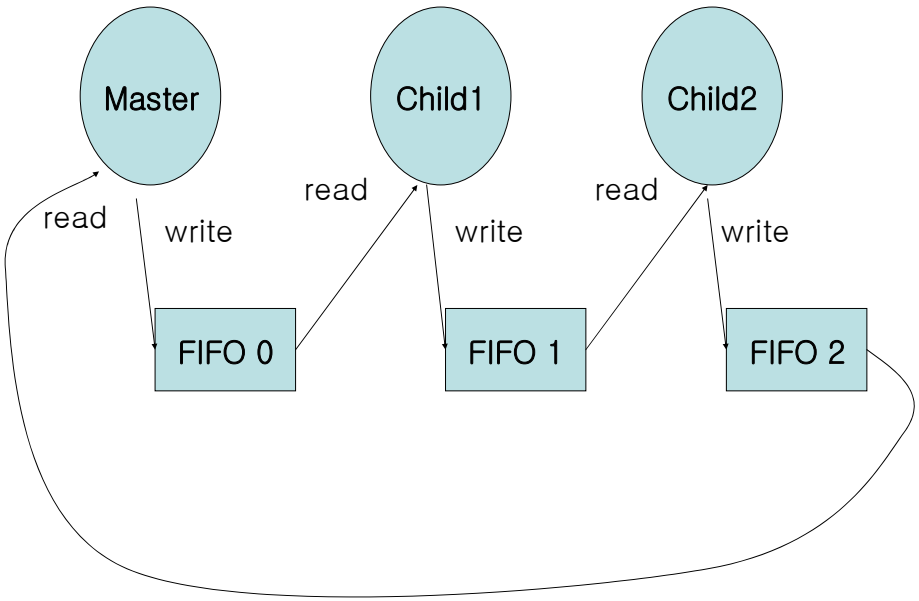

lmbench lat_ctx result

[Samsung S3C24A0(Arm926ej), 200MHz]

Clinux 2.6
nchmark-

processes	uclinux-2.6.11.6	linux-2.6.11.6	linux-2.4.20-mvista
	size=0k ovr=10.30	size=0k ovr=14.19	size=0k ovr=9.34
2	6.49	96.15	79.42
4	13.34	99.49	78.45
6	19.15	101.55	83.91
8	21.72	116.1	78.45
10	20.75	111.05	77.36
12	22.02	107.78	83.9
14	23.2	109.22	83.74
16	23.72	105.92	79.03
	size=1k ovr=18.38	size=1k ovr=19.76	size=1k ovr=15.38
2	68.91	195.78	196.79
4	125.45	217.87	196.02
6	139.46	204.04	197.2
8	147.54	205.21	196.33
10	150.56	202.95	196.31
12	151.65	208.61	195.73
14	152.41	209.37	196.62
16	153.07	207.43	196.31
	size=16k ovr=139.14	size=16k ovr=104.56	size=16k ovr=104.30
2	225.8	302.47	319.19
4	259.76	318.16	319.29
6	269.49	316.56	318.57
8	264.81	316.31	320.11
10	264.4	309.81	321.38
12	261.77	316.04	318.33
14	261.77	316.53	318.14
16	261.85	316.54	318.42

The Benchmark (4/6)

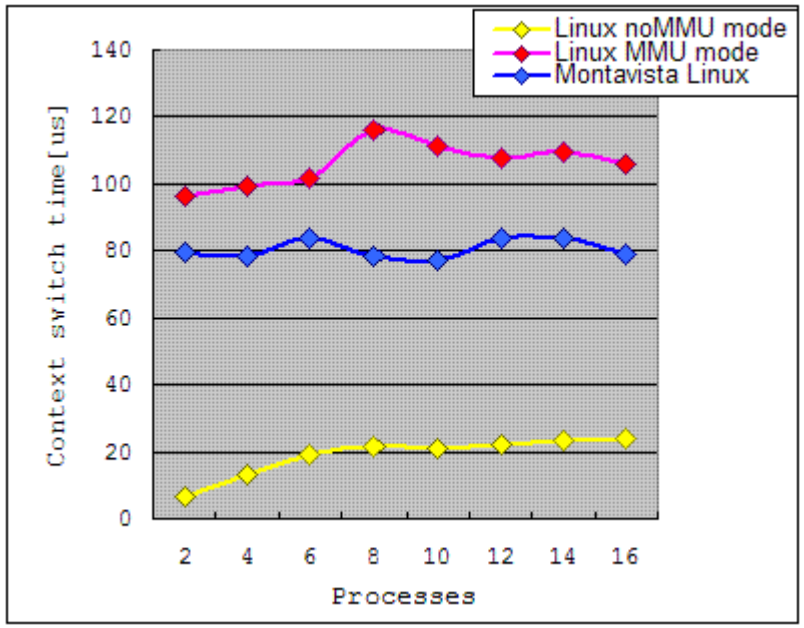

The Benchmark (5/6)

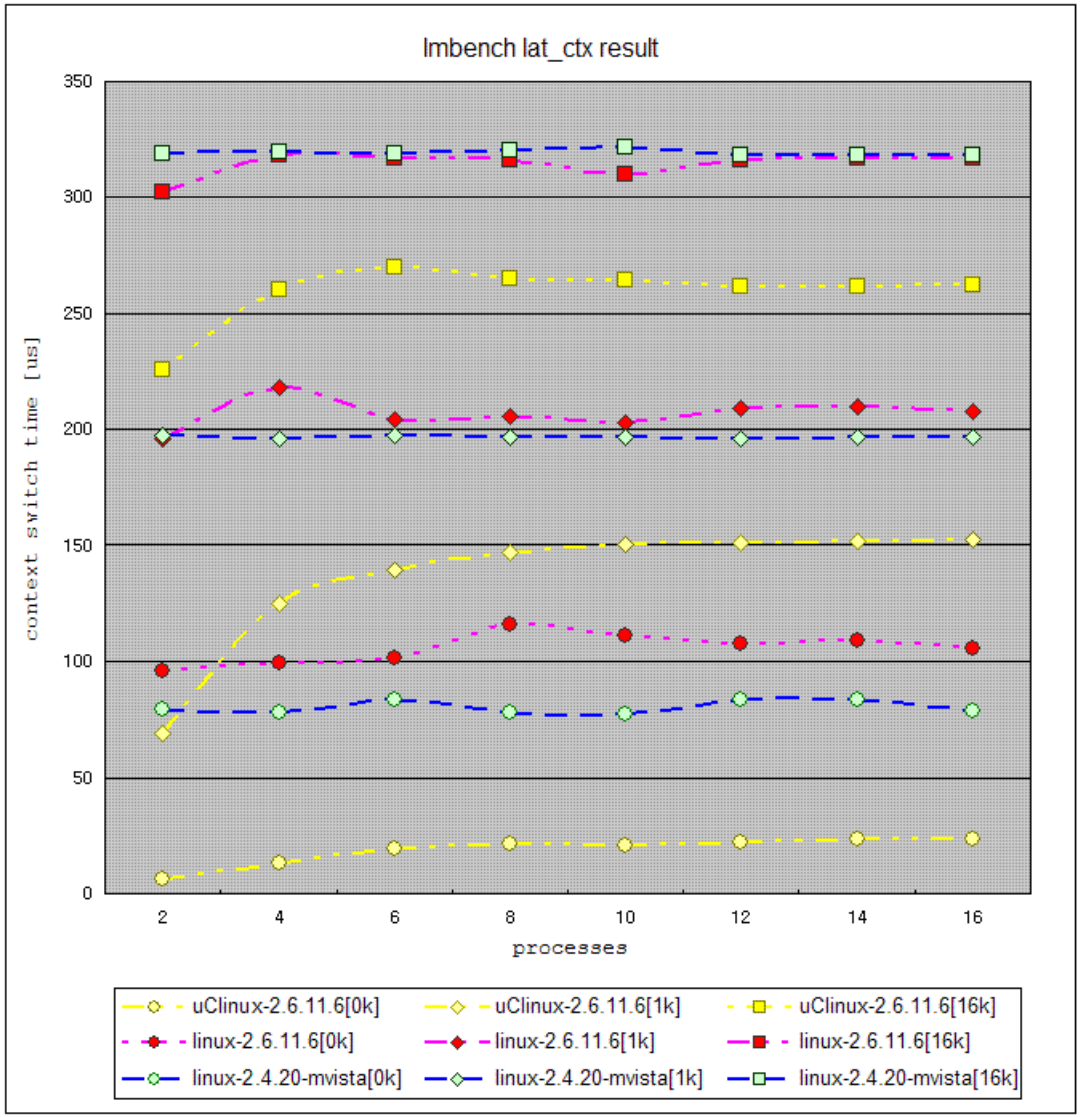

The Benchmark (6/6)

- **IPC Performance**

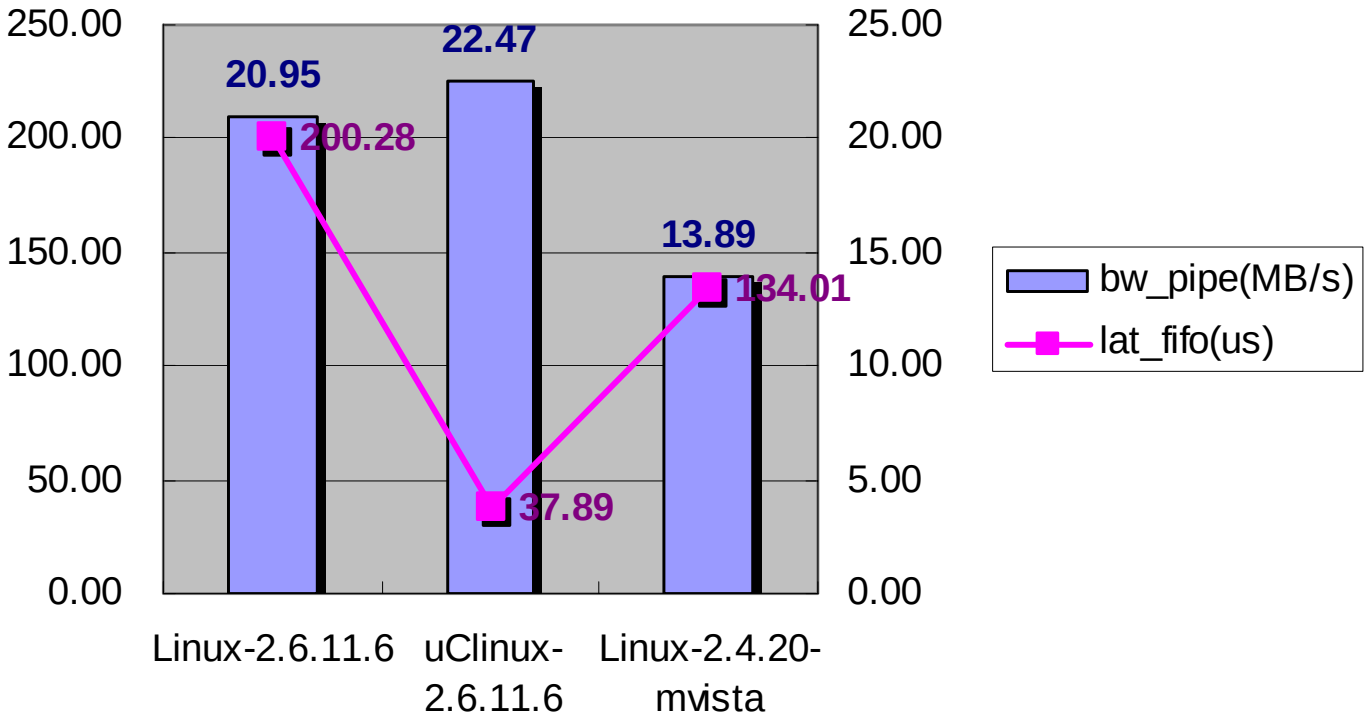

	Linux-2.6.11.6	uClinux-2.6.11.6	Linux-2.4.20-mvista
lat_fifo(us)	200.28	37.89	134.01
bw_pipe(MB/s)	(*12.58) 20.95	22.47	13.89

(*) was the result of Linux-2.6.7.

What's the Next?

- **Completion of sharing the arch with RMK**
- **Completion of V6 support**
- **MPU support**
  - Manual Memory Protection
- **Porting XScale, StrongARM**

and so on…